Sensors, Sensors Everywhere

In the near future, structures and systems will look after themselves.

OUR BODIES USUALLY LET US KNOW when they require some sort of care or treatment. Can manmade structures do the same?

"Much of our technical infrastructure is approaching, or already exceeds, its initial design life," says Chuck Farrar, leader of the Engineering Institute at Los Alamos, a research and education collaboration between Los Alamos National Laboratory and the University of California, San Diego. "We have to monitor the health of these structures because they continue to be used despite the degradation they've accumulated from their operational environments," he adds. He refers to a wide range of structures, including buildings and bridges, naval equipment and nuclear reactors, amusement park rides and aircraft, as well as large-capital scientific infrastructure items, such as particle accelerators, telescopes, and supercomputers.

Farrar wrote a textbook on structural health monitoring (SHM) and guides a cohort of early-career research staff, postdoctoral researchers, and engineering students in its methods. He believes many of the nation's infrastructure woes can be addressed with new technology that he and other members of his team are developing. The technology is designed to produce and interpret data streams from sensors that, they predict, will soon be all over the place.

If it ain't broke

Human beings are quite adept at knowing when their personal belongings need to be repaired or replaced. Some items can be judged by feel or with a simple visual inspection, as when shoe soles wear thin. Others are used until they fail and are then replaced, like a computer or a hot water heater. But there are others that can't be judged so easily by their look or feel and can't be run until they fail; they must instead be judged by time, as with the expiration date on food or medicine.

Compared to personal belongings, major infrastructure items can be much more difficult to assess. Oftentimes, they can't be run to fail because it would be either catastrophic (a bridge collapsing with people on it) or unacceptably expensive (the loss of a single machine upon which an entire production line depends). Such negative outcomes are generally prevented with regularly scheduled, or time-based, maintenance. Yet this is undesirable from a lifetime-expense perspective. It forces people, businesses, and governments to pay for inspections—or outright replacements—before they are needed. For example, when the engine oil in a car is changed every 3000 miles, even though the oil may still be usable, the owner pays for more oil changes than needed over the life of the car, to say nothing of the unnecessary environmental impact. And while additional oil changes at $40 apiece may not be too burdensome, retiring high-end hardware before its time (think combat missiles) costs considerably more.

"We can save money, gain efficiency, and improve public safety, all by shifting our culture of maintenance from time-based to condition-based with SHM," Farrar says. In that paradigm, repairs and replacements would be carried out only when they are needed. Factories wouldn't be in danger of shutting down production because one machine breaks unexpectedly (nor would backup machinery be needed as a safeguard) if the condition of the machinery were automatically monitored to provide advance notice of potential problems as they develop. Civilian and military hardware could be kept in service longer and, in some cases, relegated into less critical applications as it ages. Rental equipment could be priced according to the measured amount of wear and tear introduced by the renter. Broadly speaking, condition-based maintenance, as enabled by SHM technology, is part of a true cradle-to-grave system state awareness capability that maximizes the return on investment.

That's where the sensor proliferation comes in: sufficient data must be collected to assess each structure's condition. Unfortunately, it may be prohibitively expensive to retrofit existing structures with large numbers of sensors. (Imagine the Golden Gate Bridge needing multiple sensors on every single one of the interconnecting beam segments underlying the road surface, plus many more on the towers and cables.) However, if the sensors were incorporated into the construction of new bridges, aircraft, equipment, and so on, then their cost would amount to only a tiny fraction of the overall construction or fabrication costs. Therefore, the SHM culture shift can be expected to ramp up as major new infrastructure items are built. Indeed, this is already underway in China,

Structural health monitoring (SHM) involves deploying sensors and software to monitor a wide variety of infrastructure objects for any sign of degradation over time. Objects include buildings and bridges, power plants and industrial plants, ships and aircraft, and other large-investment equipment for transportation, entertainment, and scientific research.

where bridges, dams, offshore oil platforms, and other large infrastructure construction is accompanied by a large-scale sensing capability.

Shaky foundation—in a good way

Sensing and interpreting a wealth of structural health data is far from straightforward. Simply deploying sensors is not enough. Rather, a number of critical design decisions must be made and implemented before the SHM system is capable of delivering the desired information. For example: Which sensors should be used? How many? Where should they go? Do they need to operate in extreme temperatures or rain? Do they need to run all the time, or can they just power up periodically as needed? How often should they turn on, and how long should they stay on when they do? Will they store their measurement data until someone or something collects it, or will they transmit the data somewhere? If the latter, then how often should they transmit? And how much electrical power will be consumed by all this data acquisition, storage, and transmission? Assuming that sending repair teams to replace thousands of batteries is not an option, how will the sensors obtain the power they need, year after year?

To power the sensors without a hard-wired connection to the electrical grid, which is not always available, engineers could opt for solar energy. But while solar cells can be small and independent of the grid, their use would be restricted to locations with frequent access to direct sunlight. In a darker environment, such as the underside of a bridge, an elevator shaft, or an airplane flying at night, sunlight would not be available. What would be available, in these and a variety of other SHM settings, are frequent mechanical vibrations.

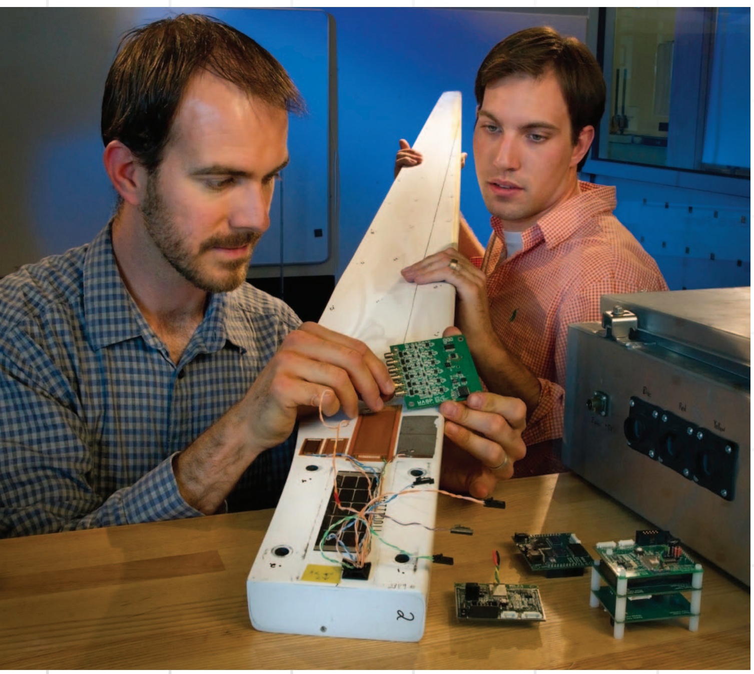

Los Alamos engineers Stuart Taylor (left) and Steve Anton fit a sensor node, capable of analyzing and transmitting sensor data, to an energy harvesting element on a wind turbine blade. Mechanical energy from the natural vibration of the blade is converted into electrical energy for use by the sensor node, allowing the sensors to continuously monitor the structural health of the blade without needing a battery or other power source.

Within certain materials, including some crystals and ceramics, mechanical stress causes electrical charge to accumulate. This property, known as piezoelectricity, has multiple practical applications. It is frequently used to make sensors that operate by converting motion into electrical signals or, in reverse, to make actuators that convert electrical inputs into motion. Los Alamos postdoctoral researchers Steve Anton and Stuart Taylor harness piezoelectricity in yet a different way, capturing and storing electrical energy from everyday vibrations to provide power in settings where light is limited.

They designed their vibration-powered sensor units with two important attributes, in addition to drawing from an energy source that's freely available in many SHM settings. One of these attributes is the energy storage element. Instead of using a rechargeable battery, which would be heavy and lose storage capacity over time, the two researchers chose to charge a supercapacitor with their captured energy. Capacitors are simple electrical storage devices consisting of two metal plates that hold equal and opposite electrical charges. "They don't hold onto a charge forever," Anton says, "but that's okay because the charge is continually replenished by more vibrations."

The other attribute Anton and Taylor built into their sensors is short-range radio transmission capability. As a result, the sensor devices are completely wireless: no wires for power coming in and no wires for sensor readings going out. Taylor explains, however, that short-range communication is a necessary limitation. "To get miles and miles of transmission capability would require more power and add a lot of weight," he says. "So sometimes we have to settle for each sensor communicating with a neighboring sensor that's no more than 50 meters away and hopping the data down the line to some data storage unit."

There are many ways to accommodate short-range communications to a local data storage unit. On an airplane, for example, wireless sensors in the wings might transmit to a central data storage unit onboard, possibly no larger than a flash drive, which could tap into the plane's internal power and connect to its radio if needed. Alternatively, the sensors could simply store the data until a separate system (or person) comes along to collect them. On an SHM-equipped bridge, a vehicle or unmanned aircraft known as a data mule could wirelessly download all the sensor data each time it drives across or flies by the bridge.

Helping hand

Once the sensors are powered and their data collected, how will that raw sensor data be converted into actionable information for damage detection, characterization, and prediction? Consider, for example, two widely used sensors: strain gauges, which measure how much a solid material is being stretched or compressed, and accelerometers, which (no surprise here) measure accelerations. Suppose a series of strain gauges and accelerometers are affixed to an airplane wing. During a flight, the strain gauges obtain a variety of measurements. Most of these readings return to normal afterward but a few remain permanently strained. The accelerometers record the motions from a wide range of forces acting in different directions with different intensities. That's the raw data.

Now, based on that data, is there any damage? Could there be damage located between two of the sensors? Is maintenance needed? Should any parts be replaced? The data alone don't answer these questions. Somehow, the pattern of measurements needs to be compared with other patterns that might indicate damage or health, even though those patterns might be specific to a particular arrangement of sensors and may not yet exist. Perhaps the patterns can be obtained through experience, with future accumulated flight hours. Or maybe they can be calculated in advance by some as-yet undiscovered formulation.

Los Alamos engineering researcher David Mascarenas, also part of the Engineering Institute, works on an unconventional solution to this data-to-decision problem. Sometimes the best way to obtain coherent, actionable information from a sensor system, he says, is to give it a helping hand—from an actual human hand.

Mascarenas took an unassuming red-knit glove and fitted it with a collection of cell phone vibrators distributed around the hand and fingers. The vibrators receive data from a set of accelerometers attached to a test object. When the test object is subjected to inputs from a shaker, various accelerometer readings cause particular parts of the glove to vibrate. In principle, the test object could be a anything from a rigid structure on a laboratory benchtop to an unmanned aircraft radioing data back whenever it climbs, banks, or changes speed. Given a little time, a person wearing the glove can learn by association what the different vibrations mean, similar to the way a driver can learn what the different vibrations of an old car mean (about to stall, bad brake rotors, etc.). From that time forward, he or she can feel what's happening to the test object, sight unseen.

"The human nervous system includes an extremely capable, extremely generalized system for interpreting sensor data, from the eyes, ears, skin, etc.," Mascarenas says. "The glove taps into that."

David Mascarenas, of the Engineering Institute at Los Alamos, models the latest fashion within the structural health monitoring community: a red glove fitted with vibrators so that he can feel what various sensor-equipped objects are doing. The wires coming from the glove go to a small unit that receives wireless signals from sensor nodes on the multi-level test structure in the room behind him where Chuck Farrar, leader of the Institute, operates a shaker connected to the structure in order to generate test signals for the glove. Over time, a person wearing the glove can learn to interpret what the structure (or any sensor data-transmitting object) is doing, based on what he or she can feel.

While computers are better at interpreting certain types of data (a barcode), humans are better with other types of data (facial recognition). Mascarenas finds that the two differing skill sets both have a role in SHM damage detection and interpretation, depending on the object and the type of damage in question. "For some applications," he says, "the combination of machine and human processing of sensor signals leads to better decision-making than either one could do alone."

Plane scan

Meanwhile, a colleague of Mascarenas and postdoctoral researcher at the Engineering Institute, Eric Flynn, is hard at work on the machine side, developing both hardware and software for automated damage recognition without any help from the red glove. Working with graduate student Greg Jarmer, Flynn designed and constructed a portable system that nondestructively probes solid surfaces for defects. The system involves a laser beam that's redirected by a motion-controlled mirror and looks like something out of a science fiction movie—a red line that sweeps across a test surface. In Flynn's experimental setup, the laser system analyzes a large metal plate, onto which corrosion damage has been introduced on the back side. A small piezoelectric vibrator shakes the plate at ultrasonic frequencies while the laser scans the undamaged side, looking for small, local changes in the resulting vibration patterns caused by the hidden corrosion.

Flynn wrote the software that his system uses to translate the laser response data into an actual damage assessment. So far, it has been proven to correctly recognize corrosion in metals and more complex damage in composite materials, as well as cracks, delamination (layers peeling apart), and holes. And unlike traditional vibration-mode analysis, it offers both high resolution and portability.

“The system could be used to scan an aircraft body between flights, using a robotic arm to move the laser all around the aircraft,” Flynn explains. “It’s not the kind of system where the airplane can test itself—we’re not quite there yet—but for slowly evolving corrosion and fatigue damage, continuous self-monitoring is much less important than comprehensive measurements. The active scan is better.”

Farrar also notes the tangible benefits. “Right now, on an unmanned aerial vehicle [UAV, or drone], there’s no way to assess damage to the wings; they are simply replaced after a certain number of flight hours. Eric’s system would save a lot of waste by preventing perfectly healthy wings from being thrown away before their time.”

Sweeping the streets

Once a set of sensors has been allocated to monitor the health of some structure or system, and a data analysis system has been established to interpret the sensor readings, the next question is how best to arrange the sensors to maximize the value of the information they collect. In addition to his laser work, Flynn has developed a technique to accomplish this optimization.

As a compelling scenario to motivate his research, he created a computer model of a generic urban environment inspired by Times Square in New York City, plus several of the surrounding blocks. He then distributed 10 (simulated) radiological detectors, capable of detecting radiation sources that may indicate the presence of a nuclear threat, such as an undetonated dirty bomb. When Flynn introduces a small radiological source somewhere within the simulated section of the city, his program calculates the range from the source to each detector, including the effects of any line-of-sight obstructions, and determines which detectors will register a signal and how strong that signal will be. If at least three detectors make solid readings (more is better), then the system should be able to narrow down where the radioactive material is located.

A technology developed at the Engineering Institute uses a scanning laser to search material surfaces for damage, such as cracking or corrosion. Having both high-resolution and portability, this system can ultimately perform a variety of important structural health monitoring functions, like checkups for unmanned aerial vehicles between flights, as depicted in this artist’s conception.

“The purpose of the simulation,” Flynn says, “is to figure out exactly where the detectors should be located to create the optimal detection probability—without knowing in advance where the source will show up or how big that source might be.” Of course, that problem can be partially avoided by simply adding more sensors, but, Flynn notes, there is always a limit to how many sensors are available for any given use. That limit could come from expense (radiological sensors requiring constructed poles for a better “view”), but it could also come from weight (sensors and wires on an aircraft or spacecraft), power requirements (multiple sensors drawing from a common battery), bandwidth (many sensors competing to transmit data over a common frequency), or data processing requirements (too much data to process quickly or cheaply). So limiting the number of detectors in the radiological simulation adds a necessary element of realism.

If circumstances allowed it, the radiological source problem might be solved with experience. Over time, many different sources would be picked up by different detectors, and the system operator might learn to tweak their locations to improve the performance. For instance, if one of the detectors never detected anything, then it could be moved around until it becomes more productive.

“Of course,” Flynn says, “time and experience are luxuries you don’t have with radiological crises.” So he set up the next best thing, simulated time and simulated experience, by implementing what’s known as a genetic algorithm: a natural-selection mechanism akin to the Darwinian process of evolving to suit the environment. The simulation initially deploys its sensors to a randomly generated set of locations and then introduces a series of simulated radiation sources, one at a time, at random locations within the section of the city under surveillance. Some of these sources represent potential threats, while others represent benign activity, such as isotopes used in routine medical procedures. With each trial source location, the performance of the detectors is evaluated. For example, how often would each detector, in conjunction with all the others, contribute valuable information? Over many different trials, the sensor network develops a lifetime of experience. Detector locations, or sets of locations, that do not help enough are like animals that are unfit for survival or reproduction; they are deleted from the "gene pool." New sensor locations are introduced to replace the ones that are deleted, based on "mutation" and "breeding" among the remaining sensor "population," in a process that repeats for each new "generation." In this way, the 3D sensor locations "evolve" to their optimal configuration.

Flynn might deservedly take pride in protecting his virtual city from nuclear disaster, but the work also has a broader impact. Optimizing the Times Square radiological alarm system is no different from optimizing the sensor placement in any surveillance problem, including SHM. Either way, it's a matter of preparing for a potentially harmful agent, whether that be a radiological weapon or a structural failure, that cannot be predicted in advance.

Sensor suicide mission

Other applications call for flexible sensor systems that can be deployed on short notice to remote and potentially hostile regions. In such cases, autonomous, sensor-packed rovers and UAVs may be needed. The researchers at the Engineering Institute, not the sort to leave any stone unturned, have designs in this arena as well.

When Mascarenas takes off his red glove, it's so he can work on his lightweight, autonomous UAV. About the size of a small bird, Mascarenas's plane is designed to extol the virtues of multitasking materials.

"For this kind of ultra-light, self-powered plane, you really can't tolerate any component that doesn't pull its own weight and then some," Mascarenas says. With that in mind, he designed the wings from a graphene oxide material that does double duty as a wing and an energy-storing capacitor. He also designed the planes to be disposable, in case they should need to be deployed into harsh environments from which they may never return.

Mascarenas's colleague Chris Stull develops software to instruct autonomous vehicles, including ground-based rovers and Mascarenas's UAVs, how to navigate within complex environments without receiving regular attention from human operators.

"We train autonomous vehicles like you might train a pet," Stull says. "Except that you probably wouldn't send your pet into a war zone."

Indeed, Stull's training program involves rewards and penalties applied to complex, dangerous tasks, such as patrolling a modern urban combat environment. "Ground-based vehicles might be rewarded for finding a place to refuel and penalized for encountering an IED [improvised explosive device]," he says. "They essentially try to maximize their score on the reward-penalty scale, all the while driving around in a pattern that will be unpredictable to the enemy."

Making it real

But even though the Engineering Institute's overall effort includes exotic, high-risk applications like these, its researchers acknowledge that real-world, sensor-system testing needs to begin with more predictable SHM environments. Indeed, sensor-based SHM is already in use in certain applications involving rotating machinery, in which operating conditions (rotation rates, temperatures) are tightly controlled, and damage scenarios (chipped gear teeth, misalignments) are well understood. On the strength of that understanding, real-world SHM applications involving rotating machinery currently in use include a variety of high-stakes platforms, such as helicopter drivetrains, Navy vessel propeller shafts, and nuclear reactor coolant pumps.

The next challenge is to employ similar technology in less controlled environments, such as aircraft subjected to a huge range of weather, ice, and turbulence conditions, possibly unmanned and behind enemy lines. With all potential applications considered, this work represents an opportunity to rebuild the infrastructure of the world to take care of itself, and the Engineering Institute's team is starting to make it happen. LDRD

—Craig Tyler

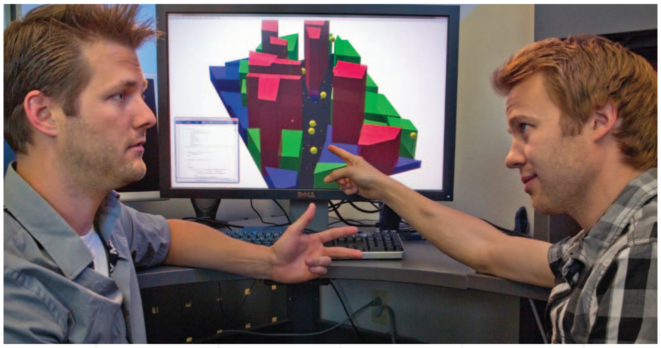

The Engineering Institute's Eric Flynn (left) and Greg Jarmer discuss the placement of radiological sources and detectors in a simulation developed by Flynn. The simulation uses a genetic algorithm that evolves throughout a series of tests until it has found an optimal detector placement. Its objective is to protect the public from potentially dangerous sources of radiation while recognizing benign ones, such as those used in routine medical scans.